

What is claimed is:

1. A cache-coherent system comprising:

a memory having a plurality of memory units;

a plurality of nodes employing a coherence protocol to maintain cache coherence of

5 the memory;

a cache within each node to temporarily store contents of the plurality of memory

units; and,

logic within each node to determine whether a cache miss relating to a memory unit

should be transmitted to one or more nodes lesser in number than the plurality of nodes,

10 based on a criteria.

2. The system of claim 1, wherein the criteria includes whether, to ultimately reach an

owning node for the memory unit, such transmission is likely to reduce total

communication traffic among the plurality of nodes and unlikely to increase latency as

compared to broadcasting the cache miss to all of the plurality of nodes.

15 3. The system of claim 1, wherein the logic within each node is to determine whether

the node is a home node for the memory unit to which the cache miss relates in

determining that transmission to the one or more nodes lesser in number than the plurality

of nodes is likely to reduce total communication traffic among the plurality of nodes and

unlikely to increase latency to ultimately reach the owning node for the memory unit.

4. The system of claim 3, wherein the one or more nodes comprises an owning node for the memory unit as stored at a directory of the home node.

5. The system of claim 1, wherein the logic within each node is to determine whether the cache of the node has stored a hint as to a potential owning node for the memory unit

5 as a result of an earlier event in determining that transmission to the one or more nodes lesser in number than the plurality of nodes is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency to ultimately reach the owning node for the memory unit.

6. The system of claim 5, wherein the event includes an invalidation of the memory unit

10 by the potential owning node.

7. The system of claim 5, wherein the one or more nodes comprises a home node of the memory unit and the potential owning node for the memory unit.

8. The system of claim 1, wherein the logic within each node is to determine whether

the memory unit relates to a predetermined memory sharing pattern encompassing the

15 one or more nodes in determining that transmission to the one or more nodes lesser in number than the plurality of nodes is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency to ultimately reach the owning node for the memory unit.

9. A method comprising:

determining at a first node whether a cache miss relating to a memory unit of a shared memory system of a plurality of nodes including the first node and employing a coherence protocol should be selectively broadcast to one or more nodes lesser in number than the plurality of nodes based on a criteria;

5 in response to determining that the cache miss should be selectively broadcast to the one or more nodes, selectively broadcasting the cache miss by the first node to the one or more nodes.

10. The method of claim 9, further comprising, in response to determining that the cache

10 miss should not be selectively broadcast to the one or more nodes, broadcasting the cache miss by the first node to all of the plurality of nodes.

11. The method of claim 9, wherein the criteria includes whether selective broadcasting is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to just broadcasting the cache miss to all of the plurality of
15 nodes to reach an owning node for the memory unit.

12. The method of claim 9, wherein determining whether the cache miss should be selectively broadcast to the one or more nodes comprises determining whether the first node is a home node for the memory unit, such that selectively broadcasting the cache miss to the one or more nodes comprises selectively broadcasting the cache miss to one
20 node of the plurality of nodes as an owning node for the memory unit as stored at a directory of the first node as the home node for the memory unit.

13. The method of claim 9, wherein determining whether the cache miss should be selectively broadcast to the one or more nodes comprises determining whether the first node has a pre-stored hint as to a potential owning node for the memory unit, such that selectively broadcasting the cache miss to the one or more nodes comprises selectively broadcasting the cache miss both to a home node of the memory unit and to the potential owning node for the memory unit.

5
14. The method of claim 9, wherein determining whether the cache miss should be selectively broadcast to the one or more nodes comprises determining whether the memory unit relates to a predetermined memory sharing pattern encompassing the one or 10 more nodes, such that selectively broadcasting the cache miss to the one or more nodes comprises selectively broadcasting the cache miss to the one or more nodes.

15. A method comprising:
determining at a first node whether a cache miss relating to a memory unit of a shared memory system of a plurality of nodes including the first node and employing a 15 coherence protocol should be selectively broadcast to one or more nodes lesser in number than the plurality of nodes, based on whether selective broadcasting is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to just broadcasting the cache miss to all of the plurality of nodes to reach an owning node for the memory unit; and,
20 in response to determining that the cache miss should be selectively broadcast to the one or more nodes, selectively broadcasting the cache miss by the first node to the one or more nodes.

16. A method comprising:

determining at a first node whether a cache miss relating to a memory unit of a shared memory system of a plurality of nodes including the first node should be selectively broadcast to one or more other nodes of the plurality of nodes, based on whether the first
5 node is a home node for the memory unit or whether the first node has a pre-stored hint as to a potential owning node for the memory unit;

in response to determining that the cache miss should be selectively broadcast to the one or more other nodes, selectively broadcasting the cache miss by the first node to the one or more other nodes;

10 otherwise, determining at the first node whether the memory unit relates to a predetermined memory sharing pattern encompassing a sub-plurality of the plurality of nodes smaller in number than the plurality of nodes; and,

in response to determining that the memory unit relates to the predetermined memory sharing pattern, selectively broadcasting the cache miss by the first node to the sub-plurality of the plurality of nodes.

17. A node of a system having a plurality of nodes comprising:

local memory for which the node is a home node and that is shared among the plurality of nodes;

a directory to track which of the plurality of nodes has cached or modified the local
20 memory of the node;

a cache to temporarily store contents of the local memory and memories of other ones of the plurality of nodes; and,

logic to determine whether a cache miss relating to a local memory should be

transmitted to one or more nodes lesser in number than the plurality of nodes based on whether, to ultimately reach an owning node for the local memory, such transmission is likely to reduce total communication traffic among the plurality of nodes and unlikely to increase latency as compared to broadcasting the cache miss to all of the plurality of
5 nodes.

18. An article of manufacture comprising:

a computer-readable medium; and,

means in the medium for selectively broadcasting a cache miss relating to a memory unit of a shared memory system of a plurality of nodes employing a coherence protocol
10 to one or more nodes lesser in number than all the plurality of nodes of the shared memory system, based on a criteria.

19. The article of claim 18, wherein the means is for selectively broadcasting the cache miss to an owning node for the memory unit where an originating node of the cache miss is a home node for the memory unit.

15 20. The article of claim 18, wherein the means is for selectively broadcasting the cache miss to a home node for the memory unit and a potential owning node for the memory unit where an originating node of the cache miss has at a cache thereof a pre-stored hint as to the potential owning node as a sending node of an earlier received invalidation of the memory unit.

21. The article of claim 18, wherein the means is for selectively broadcasting the cache miss to a sub-plurality of the plurality of nodes smaller in number than the plurality of nodes where the memory unit relates to a predetermined memory sharing pattern encompassing the sub-plurality of the plurality of nodes.

5